

A STRATEGY FOR ASSESSING CLOSURE

PETER MURPHY

ABSTRACT

This paper looks at an argument strategy for assessing the epistemic closure principle. This is the principle that says knowledge is closed under known entailment; or (roughly) if *S* knows *p* and *S* knows that *p* entails *q*, then *S* knows that *q*. The strategy in question looks to the individual conditions on knowledge to see if they are closed. According to one conjecture, if all the individual conditions are closed, then so too is knowledge. I give a deductive argument for this conjecture. According to a second conjecture, if one (or more) condition is not closed, then neither is knowledge. I give an inductive argument for this conjecture. In sum, I defend the strategy by defending the claim that knowledge is closed if, and only if, all the conditions on knowledge are closed. After making my case, I look at what this means for the debate over whether knowledge is closed.

According to the epistemic closure principle (hereafter, Closure), if someone knows that *P*, knows that *P* entails *Q*, goes on to infer *Q*, and in this way bases their belief that *Q* on their belief that *P* and their belief that *P* entails *Q*, then they know that *Q*. Similar principles cover proposed conditions on knowledge. Call the conditions on knowledge, whatever they might be, *k-conditions*. A *k-condition*, *C*, is closed if and only if the following is true: if *S* meets *C* with respect to *P* and *S* meets *C* with respect to *P* entails *Q*, and *S* goes on to infer from these to believe *Q* thereby basing her belief that *Q* on these other beliefs, then *S* meets *C* with respect to *Q*.¹

On one view about the relationship between knowledge and its conditions, Closure is true just in case all *k-conditions* are closed. Call this the Equivalence Claim. It factors into two claims. One is

(AC) If Closure is true then all k-conditions are closed.

This claim might be of value to someone who wants to argue, in a modus tollens fashion, against Closure. I will call it ‘AC’ for *anti-Closure*. The other claim might be of value to someone who wants to argue for Closure. I will call it ‘PC’ for *Pro-Closure*. It says

(PC) If all k-conditions are closed, then Closure is true.

The aim of this paper is to assess these claims and determine what role they might play in the debate over Closure. To anticipate, I will argue for both AC and PC, therein arguing for Equivalence. Then I will argue that Equivalence has some important implications for the debate over Closure.

In Section 1, I will provide reasons for thinking that Closure is important. In Section 2, I will give a proof for PC, one that makes no assumptions about what the correct k-conditions are. In Section 3, I will argue that AC cannot be established without making such assumptions. Still, over the course of Sections 4 and 5, I make an inductive case for AC, one that takes into account what the correct k-conditions might be. I end in Section 6 by highlighting some important implications that Equivalence has for the contemporary debate over Closure.²

1. WHY CLOSURE MATTERS

There are a number of reasons for thinking that Closure is important and merits investigation. Let me review two. The first concerns the role that Closure plays in a key argument for skepticism. The argument has three premises. Premise

1: If I know that (i) I have hands and I know that (ii) if I have hands then I am not a handless BIV who is being electrochemically stimulated so that it merely appears to me that I have hands, then I know that I am not such a BIV. Premise 2: I do not know that I am not such a BIV. Premise 3: I know the conditional at (ii). Conclusion: I do not know that I have hands. This argument is valid. Moreover, Premise 3 is difficult to dispute. This means that if Premises 1 and 2 are correct, skepticism is correct. But Premise 1 is true if Closure is true. For this reason, Closure is crucial to the assessment of this important argument.

Closure is also important for determining how knowledge behaves in various patterns of inference. Consider two important principles that require the truth of Closure. One principle has it that knowing Q is something that one must *antecedently accomplish* if one is going to know P , know P entails Q , and be able to infer from these to Q .³ In other words, knowing P , knowing P entails Q , and being able to infer from these beliefs to Q requires first having independent knowledge of Q . The other principle, known as ‘the Transmission Principle’, is weaker. It does not require first having independent knowledge of Q . Instead, it says that if S knows that P , S knows that P entails Q , and S appropriately infers from these to Q , S can thereby come to know Q *for the very first time*.⁴ Though both of these principles are stronger than Closure, each requires the truth of Closure.⁵ Hence, Closure must be defensible if either of these principles is to be defended.

2. A PROOF FOR PC

There is a straightforward proof for PC, the claim that if all k-conditions are closed, then Closure is true. For the sake of conditional proof, assume that all k-conditions are closed. Does Closure follow? Since Closure has the form of a conditional statement, assume its antecedent. It is the conjunction of four statements: S knows that P , S knows that P entails Q , S infers Q , and S therein bases the belief that Q on the belief that P and the belief that P entails Q . The first conjunct implies that S meets each of the k-conditions with respect to P . The second conjunct implies that S meets each of the k-conditions with respect to P entails Q . Next isolate each k-condition. From the claim that S meets the isolated condition with respect to P , the claim that S meets it with respect to P entails Q , plus the third and fourth conjuncts in the antecedent of Closure, as well as the lead assumption that the isolated condition is closed, it follows that S meets the condition with respect to Q . Repeating this reasoning for each k-condition allows us to show that S meets all the k-conditions with respect to Q . Therefore, S knows that Q . This proves PC.

3. AC AND THE CONTAINMENT PRINCIPLE

Things get more complicated when we turn to AC, the claim that if one (or more) of the k-conditions is not closed, then knowledge is not closed.⁶ It is natural to try to argue for this claim by conditional proof. Call the k-condition that is not closed, C. Then begin with the assumption that C is not closed. Assume, that is,

that (1) S meets C with respect to P , (2) S meets C with respect to P entails Q , and (3) S does not meet C with respect to Q .⁷ From these claims, does it follow that Closure is false? That is, does the conjunction of these claims follow: (4) S knows that P , (5) S knows that P entails Q , and (6) S does not know that Q ?

(6) follows; it follows from (3). However, there is no way to derive either (4) or (5), since (1)-(3) fail to imply anything about whether S meets k-conditions other than C with respect to P , as well as with respect to P entails Q . So, from the sole fact that one k-condition is not closed, it does not follow that Closure is false.

However, there is a supplemental claim that, when added to (1)-(3), yields the falsity of Closure. To see what this claim is, notice that what we need is that S meets all the k-conditions, besides C, with respect to both P and P entails Q . If S does this, S will know P , and S will know P entails Q . Add these last two claims to the claim that S fails to know Q , and it follows that Closure is false. This means that if S's failure to meet the non-closed condition, C, with respect to Q does not entail that S fails to meet some other k-condition (besides C) with respect to either P or P entails Q , there will be room for Closure to be false. As I will put it, as long as S's failure to meet C with respect to Q is *contained* in the sense that it does not spoil S's standing with respect to either P or P entails Q , Closure can fail. Call the claim that this is so, the Containment Principle:

(CP) S's failure to meet a non-closed k-condition, C, with respect to Q does not entail that S fails to meet any other k-condition, D, with respect to either P or P entails Q .

If the Containment Principle is true, then C's failure to be closed makes it possible for the relevant instance of Closure to be false. But since Closure is the kind of claim that is true if and only if it is necessarily true, the possibility that it is false entails that it is false.

In a moment, I will give an inductive argument for the Containment Principle. But first I want to consider a bolder argument for the Containment Principle. Anthony Brueckner offers a short deductive argument for this principle. He argues that it follows from the fact that k-conditions need to be independent of one another. Brueckner points out that if the Containment Principle were false then a failure to meet C with respect to *Q* would entail a failure to meet D with respect to either *P* or *P* entails *Q*. This is equivalent to this disjunction: either meeting D with respect to *P* entails meeting C with respect to *Q*, or meeting D with respect to *P* entails *Q* entails meeting C with respect to *Q*. Focusing on the former, Brueckner asserts that “[i]f the conditions for knowledge are independent of each other, then the satisfaction of D by S with respect to *P* will not entail the satisfaction of C by S with respect to *Q*.”⁸ Since the conditions for knowledge are independent of each other, there is no such entailment. So, the Containment Principle is true.

Brueckner's case turns on a claim about a particular kind of epistemic independence. What should we make of this kind of epistemic independence? Should we require that it hold? Notice, first, that the kind in question involves independence vis-à-vis two distinct k-conditions (C and D), with respect to two

distinct propositions (P and Q). Brueckner provides no supporting reasons for requiring this independence. Moreover, methodological niceties alone do not allow us to ban the relevant dependency claims. The claims that would be banned are of this form: S meets D with respect to P only if S meets C with respect to Q . Consider the two most obvious sources of methodological strictures: the need to avoid regresses and the need to keep out redundancies. Dependency claims of this form do not introduce either a regress or a redundancy into the analysis of S knows P . A regress need not ensue, since the dependency claim does not entail that S meets C with respect to Q only if S meets some k -condition (C , D , or any other) with respect to some third proposition, R . Redundancy is not introduced, since a distinct condition, C , is imposed with respect to a distinct second proposition, Q , and not with respect to P .

These dependency claims, involving two distinct k -conditions and two distinct propositions, contrast with other kinds of epistemic dependency claims that are objectionable. Consider, first, claims of this form: S meets D with respect to P only if S meets C with respect to P . This claim is redundant: all that has to be mentioned when citing the conditions that S must meet to know P is that S meets D ; there is no need to add that S must meet C . Since analyses should be free of redundancies, we can reasonably reject claims like this (at least, as long as the claims show up in the analysis of S knows P). Consider, second, claims of this form: S meets D with respect to P only if S meets D with respect to Q . A claim like this may trigger a regress. It will do this if the reason for imposing the

condition is perfectly general. If it is perfectly general, it will apply to S's meeting D with respect to *Q*. Consequently, S will meet D with respect to *Q* only if S meets D with respect to *R*, and a regress will be triggered. To avoid this kind of regress, we need to ban this form of epistemic dependence.

The kind of dependency claim that Brueckner wants to rule out is quite different. It introduces no redundancy and it does not trigger a regress. So there are no clear reasons to reject it. Still, I think there are good reasons to embrace the Containment Principle. But these reasons do not stem from any methodological stricture. Instead, they stem from the nature of k-conditions and how they interact with one another. So to support the Containment Principle, I am going to investigate various k-conditions to see, first, which ones are closed. Then I will determine whether the non-closed k-conditions are contained vis-à-vis other k-conditions. The first task will be taken up in the next section; the second task in the following section. After surveying a numerous and diverse pool of k-conditions and showing that the Containment Principle stands up, I will be in a position to claim that AC has been inductively supported.⁹

4. A SURVEY OF K-CONDITIONS: CLOSED AND NOT CLOSED

I will begin with the non-epistemic k-conditions, truth and belief; then I will look at some familiar externalist-style conditions, a representative internalist condition, and a no defeater condition.

Truth is obviously closed: if it is true that P and it is true that P entails Q , then it is true that Q . Belief is more complicated. On most theories of belief, belief is not closed: even if S believes that P and S believes that P entails Q , S may not put these beliefs together to form the belief that Q . People are often surprised at what their own beliefs entail. Recall, though, that the formulation of Closure that I am working with also includes in its antecedent a clause which says that S goes on to believe Q and that S does this on the basis of S 's beliefs that P and P entails Q . What motivates including these clauses is the thought that the preceding considerations about belief are no threat to the spirit of Closure. The spirit of Closure concerns *epistemic* credentials: if someone has the epistemic credentials needed to know P , has the same credentials with respect to P entails Q , and infers and thereby comes to believe Q , must one have those credentials with respect to Q ? The credentials in question are epistemic in nature – they are over and above the truth and belief conditions. With this in mind, let's look at some epistemic credentials. I will examine six kinds of credentials; I will argue that five of them are not closed; I will also argue against a strategy that someone might use to show that a condition is closed.

Sensitivity is not Closed: Dretske's Zebra

According to the sensitivity condition, S knows that P only if S 's belief that P is sensitive. For that belief to be sensitive, the following counterfactual must be true: if P were false, S would not believe that P . Sensitivity, as often noted, is not

closed. It is possible for S to sensitively believe that *P*, sensitively believe that *P* entails *Q*, from these beliefs deduce (and thereby believe) *Q*, yet fail to sensitively believe that *Q*.

Here is a familiar case that will also serve us later. Fred sees that there is a zebra in front of him. Believing that if there is a zebra in front of him then there is not a cleverly disguised mule in front of him, Fred infers that there is not a cleverly disguised mule in front of him. Fred's belief that there is a zebra in front of him is sensitive: if there were not a zebra in front of him, then he would not believe there was. His belief in the entailment is also sensitive.¹⁰ However, his belief that there is not a cleverly disguised mule in front of him is not sensitive: for, if there were a cleverly disguised mule there, Fred would nonetheless believe that there was not a cleverly disguised mule there.

The Causal Condition is not Closed: Dretske's Zebra Again

Next is a simple version of a causal condition, one that says S knows that *P* only if S's belief that *P* is suitably caused by *P*. This condition has to be supplemented with some way of figuring out when causal relations are present. Once we have such a theory in place, we can ask: if one meets the causal condition with respect to *P* and one does the same with respect to *P* entails *Q*, and one then deductively infers *Q*, must one also meet the causal condition with respect to *Q*? Dretske's Zebra case doubles to show that the causal condition is not closed. Fred's belief that there is a zebra in front of him is suitably caused by the fact that

there is a zebra in front of him. To support this claim, we need some reliable way of determining if a causal relation is present. Consider, then, *c* causes *e* if and only if: modulo identifiable kinds of counterexamples, if *c* had not occurred, *e* would not have occurred. Identifiable kinds of counterexamples include, most notably, early, late, and trumping preemption cases, as well as prevention cases.¹¹ With this in mind, what should we say about Fred's belief that there is a zebra in front of him? As we saw earlier, if there were not a zebra in front of Fred, he would not have believed there was a zebra in front of him. Moreover, since this case is not like any of the familiar kinds of counterexamples to the counterfactual theory of causation, it is highly plausible that the causal relation is present.

What about Fred's belief that if there is a zebra in front of him then there is not a cleverly disguised mule there? It is well known that defenders of the causal condition have had a difficult time dealing with beliefs like this. The difficulty arises because conditional facts are arguably abstract, and therefore not fit to be causes. For the sake of determining whether the causal condition is closed, let's suppose that this difficulty can be overcome. That gets us to Fred's belief that there is not a cleverly disguised mule in front of him. As noted earlier, if there were a cleverly disguised mule in front of Fred, he would nonetheless believe that there was not such a mule there. This suggests that the needed causal relation is not in place. Nor can this case be assimilated to any of the counterexamples to the simple counterfactual theory. So it is reasonable to think that what makes this

belief true is not the cause of this belief. The causal condition on knowledge is, therefore, not closed.

Safety is not Closed: Lying Larry

There is another important counterfactual condition, safety. It is not closed either.¹² According to this condition, S knows that *P* requires the truth of a different counterfactual, namely this one: if S were to believe *P*, then *P* would be true. That safety is not closed is illustrated by the following case. Among my casual acquaintances is Larry. I have no reason to doubt what Larry says. One day he reports to me that he is married to Pat. I believe Larry and I also believe that if Larry is married to Pat then he is married to someone; so, I go on to infer that Larry is married to someone. My belief that Larry is married to Pat is safe: in the nearest (non-actual) worlds in which I believe it, it is true. My belief in the entailment is also safe: it too holds in the nearest worlds in which I believe it. What about my belief that Larry is married to someone? Well suppose that Larry's proposal to Pat had a very low chance of succeeding. Moreover, suppose that Larry had a backup plan: had Pat turned him down, he would have deceived me into thinking that he was married. However, he would never have lied about being married to *Pat*. Instead, he would have told me that he was married to someone else. Because his proposal to Pat was a longshot and he was disposed to lie in this way, in the nearest (non-actual) world in which I believe Larry is married to someone, he is not married. So while my belief that Larry is married to Pat is safe,

and my belief in the entailment is safe, my belief that Larry is married to someone is not safe. So safety is not closed.

Reliability is not Closed: The Logic Student

Consider next the process reliabilism condition. It says that S knows that P only if the belief forming process responsible for S's belief that P produces a sufficiently high ratio of true to false beliefs. Being reliably produced is not closed either. That the process responsible for the belief that P and the process responsible for the belief that P entails Q are both reliable, does not guarantee that the process responsible for the belief that Q is reliable.¹³ Consider the following case. A student believes P and believes P entails Q . Both of these beliefs were reliably produced. He then goes on to infer Q from these beliefs. However, he does this by employing a process that yields the belief that Q whenever Q shows up in any proof with two, or more, premises.¹⁴ Since that is not a reliable process, the student's belief that Q is not reliably formed.¹⁵ Still, the belief that P and the belief that P entails Q were both reliably produced. So being produced by a reliable belief forming process is not closed.

The Standard Internalist Condition is not Closed: The Math Student

How about a representative internalist condition, understood as a condition that requires knowers to have an introspectively accessible psychological state that indicates the truth of the proposition believed? A condition of this sort is not

closed, either. Consider the following case. A math student has an a priori insight into an entailment between two mathematical claims. She also has an a priori insight into the antecedent claim. She then goes on to believe the entailed proposition, and she does so on the basis of deducing it from her other two beliefs. These facts alone do not require that she have an a priori insight or any other introspectively available indication that the deduced proposition is true.

A Resistance Strategy

At this point, it is worth considering a resistance strategy that defenders of some of these conditions might use. Since similar considerations apply to each deployment of the strategy, I will confine my discussion to an internalist deployment. In the math student case, the internalist might claim that the internalist justifier that the student has for her belief in the antecedent, the internalist justifier for her belief in the entailment, along with her basing the deduced belief on these beliefs together *suffice to constitute* an introspectible state that indicates the truth of the deduced proposition. No separate, independent justifier is needed for the belief in the deduced proposition.

This sort of strategy hinges on the claim that a k-condition can be closed in virtue of just the following facts: someone meets the condition with respect to P , they do the same with respect to P entails Q , and their belief that Q meets the basing requirement. On the suggested strategy, these things are enough to meet the internalist condition with respect to Q . This contrasts with a tougher standard,

according to which a k-condition is closed only if what makes one satisfy the condition with respect to Q is distinct from, in addition to, and not constituted by, whatever makes one satisfy it with respect to P , whatever makes one satisfy it with respect to P entails Q , and whatever makes one meet the basing requirement. As *The Math Student* shows, when the tougher standard is imposed, the representative internalist condition comes out non-closed.

We need to adjudicate between the two standards. The weaker standard, I submit, is too weak. To see that we need more than the weaker standard, consider a proponent of the sensitivity condition who is embarrassed by the thought that sensitivity is not closed. Suppose he attempts to avoid the embarrassment by claiming that (1) having a sensitive belief that P , (2) having a sensitive belief that P entails Q , and (3) basing the belief that Q on these beliefs together constitute (4) having a sensitive belief that Q . This should only be seen as a desperate measure. (1)-(3) simply do not constitute having a sensitive belief that Q . To have a sensitive belief that Q requires that one not believe Q in the nearest not- Q worlds – plain and simple. But, if the weaker standard were the correct standard, the foregoing would successfully show that sensitivity is closed. Since it shows no such thing, the weaker standard has to be rejected.

Someone might claim that the weaker standard is appropriate for determining that the internalist condition is closed, but that it is not appropriate for determining whether sensitivity is closed. But this sort of differential treatment would need to be motivated. We would need to know why one standard is

appropriate for determining whether some k-conditions are closed, while another standard is appropriate for determining whether others (like the sensitivity condition) are closed. Until such rationales are supplied, we should opt for the tougher standard; and on that standard, our representative internalist condition comes out non-closed.

The No Defeater Condition is Closed

Let's look at one final k-condition, the no defeater condition.¹⁶ To say that there is no defeater for S's belief that *P* based on *e* is to say that there is no undefeated defeater for S's belief that *P*. The notion of an undefeated defeater is typically spelled out in two steps. First, a proposition *d* defeats one's knowledge that *P* based on evidence *e* if and only if *d* is true and the conjunction of *d* and *e* would no longer justify S in believing *P*. Second, *d* does this in a way that is itself undefeated if, in addition, there is no further proposition *f* such that *f* is true and the conjunction of *d*, *e*, and *f* would once again justify S in believing *P*.

Is the following state of affairs possible? There is no defeater for one's justified belief that *P*, there is no defeater for one's justified belief that *P* entails *Q*, one deduces and thereby believes *Q*, and yet there is a defeater for one's justified belief that *Q*? The answer, albeit a tentative one, is 'no.' For consider what follows from there being a defeater for one's justified belief that *Q*. Suppose, for example, that a mathematics student derives *Q* from *P* and *P* entails *Q*. The student is then exposed to a proof for *not-Q*; it defeats her justification for believing *Q*. As a

result, the student is in the epistemic situation of someone who faces a reductio ad absurdum of her beliefs that P and P entails Q . This would defeat her justification for believing at least one of P and P entails Q . This suggests that the no defeater condition is closed.¹⁷

Let's summarize. Some k-conditions are closed: the truth condition and the undefeated justification condition fit here. Another k-condition, though not closed, can be finessed to come out closed without spoiling the epistemic spirit of closure: this is where the belief condition falls. That leaves the k-conditions that are not closed, and cannot be finessed. Five conditions fit into this category: the sensitivity condition, the safety condition, the reliabilist condition, the causal condition, and our representative internalist condition.

5. HOW THE K-CONDITIONS INTERACT

Let's now run these results through the Containment Principle. Again, it says

(CP) S's failure to meet a non-closed k-condition, C, with respect to Q does not entail that S fails to meet any other k-condition, D, with respect to either P or P entails Q .

Since closed k-conditions are not candidates for C, we can set aside the undefeated justification condition and focus on the five k-conditions that are not closed. I will

use the same cases to show that none of these conditions count against the Containment Principle.

Recall Dretske's Zebra Case. It was used to show that sensitivity is not closed. Fred's belief that there is not a cleverly disguised mule in front of him – call this, his *m*-belief – is not sensitive. He arrived at this belief by deduction from his belief that there is a zebra in front of him – call this, his *z*-belief – and his belief in the relevant entailment. However, the fact that his *m*-belief is not sensitive does not take anything away from his *z*-belief or his belief in the entailment. For Fred need not fail to meet any *k*-condition with respect to either *z* or the entailment. It is consistent with the case as described that with respect to both *z* and the entailment, Fred meets the truth condition, the safety condition, the reliabilist condition, the causal condition, and the representative internalist condition. In fact, the case that Dretske seems to want us to consider is a mundane case in which Fred meets all these other conditions with respect to both *z* and the entailment.

Things are not as obvious when D is occupied by the no defeater condition.¹⁸ From the fact that Fred's *m*-belief is not sensitive, does it follow that either the justification for his *z*-belief is defeated or the justification for his belief that *z* entails *m* is defeated? If the fact that Fred's *m*-belief is insensitive is added to the evidence that he has for *z* and is also added to his evidence for *z* entails *m*, does the correct theory of justification have it that he is no longer justified in holding one of these last beliefs? Maybe it does, since Fred will have been

informed that a close entailment of what he believes is something that he fails to sensitively believe.

We should be cautious, though. For one, it does not look like defeat ensues if sensitivity is only a *candidate* k-condition and not among the *actual* k-conditions. For if it is only a candidate k-condition, the defeating status of *Fred's belief that m is insensitive* will itself be defeated by the further addition of *sensitivity is not an actual condition on knowledge*. On the other hand, if it turns out that sensitivity is among the actual k-conditions, the defeater could be strengthened so that it reads: *Fred's belief that m is insensitive & sensitivity is a condition on knowledge*. Still this might not be enough to constitute a defeater for either Fred's z-belief or his belief that z *entails* m. For even when Fred comes to believe that his belief in the inferred proposition is insensitive, it does not follow (as we have seen) that either of these last two beliefs is insensitive, mistaken, or in any other way epistemically defective. This suggests that sensitivity is contained vis-a-vis undefeated justification.

Our next non-closed condition is the causal condition. We saw that it too exhibits closure failure in the zebra case. I will argue that it too respects the Containment Principle. Notice that in this case too, the fact that Fred's m-belief is not appropriately caused need not take anything away from either his z-belief or his belief that z *implies* m. As we just saw, the last two beliefs are nonetheless sensitive, safe, reliably produced, and conform to the simple internalist condition. Moreover, the same considerations apply when D is occupied by the no defeater

condition. Those considerations run as follows. Suppose, first, that the fact that Fred's *m*-belief fails to be appropriately caused is added both to his evidence for *z* and his evidence for *z* entails *m*. Either the causal condition is among the actual *k*-conditions, or it is not. If it is not, then the fact that it is not neutralizes any defeating power that the added evidence might have had; so, defeat does not ensue. If it is among the actual *k*-conditions, defeat may still not ensue, since even with the addition, it does not follow that Fred's *m*-belief or his belief that *z* entails *m* is, in any independent way, epistemically defective.

Safety also respects Containment. Recall the case involving Larry. My failure to have a safe belief that Larry is married to someone does not entail that my belief that Larry is married to Pat fails to meet any other *k*-condition, nor does it follow that my belief in the relevant entailment (i.e. if Larry is married to Pat then he is married to someone) fails to meet any other *k*-condition. We can imagine that both of these beliefs meet all of the other *k*-conditions: they are true, sensitive,¹⁹ reliably formed, and they meet the causal condition and the internalist condition.

Once again, things are more complicated when we turn to the undefeated justification condition. Here, the considerations that applied in the discussions of sensitivity and safety replay themselves. If safety is not among the actual *k*-conditions, this fact can be added to *my belief that Larry is married to someone* is *unsafe* with the result that defeat does not seem to ensue. If safety is a genuine *k*-condition, then when *safety is a genuine k-condition* and *my belief that Larry is*

married to someone is unsafe is added to my evidence for *Larry is married to Pat* and to my evidence for the entailment, it is still not obvious that either of the latter two beliefs will no longer be justified. Again, the worry is that the belief that *Larry is married to someone* is unsafe does not entail that either of the latter beliefs is in any way epistemically defective.

Similar things apply regarding *The Logic Student* case that was used to illustrate that being reliably produced is not closed. That case can be filled in so that the logic student meets all the k-conditions with respect to both of the premises from which he begins. Both beliefs, we can suppose, are true, sensitive, safe, appropriately caused, justified by internalist standards, and indefeasibly justified. So too with *The Math Student* case in which internalist justification failed to be closed. The math student's starting beliefs, we can suppose, meet all of our k-conditions. Once again, the failure to meet the k-condition with respect to the proposition *Q* seems well contained. It does not entail that the belief in the *P*-proposition or the belief in *P* entails *Q* fails to meet any other condition.

So we can conclude, albeit inductively, that if any of the non-closed conditions surveyed here shows up among the actual k-conditions, they will respect the Containment Principle and consequently Closure will be false.²⁰

6. IMPLICATIONS FOR THE CLOSURE DEBATE

So far we have a deductive case for PC and an inductive case for AC. Together these yield the Equivalence Claim: Closure is true just in case all k-

conditions are closed. What effect should the adoption of Equivalence have on the debate over Closure, a debate that goes back to Fred Dretske's important papers from the early 1970s? Since then most epistemologists have lined-up in favor of Closure. Still, those who have opposed it are a formidable bunch, among them Dretske himself, Robert Nozick, Alvin Goldman, Colin McGinn, and Robert Audi.²¹ What implications does my defense of Equivalence have for this debate?

Let's begin with some straightforward implications. Equivalence entails that two views are false. One view combines an analysis of knowledge on which all k-conditions are closed with the denial of Closure. This view runs contrary to PC; for this reason, it is false. As far as I can tell, this view is not endorsed by anyone. However, another view that Equivalence rules out has been defended. Ted Warfield extends a helping hand to the view that combines an analysis on which one (or more) k-condition is not closed with the view that Closure is correct.²² This view runs contrary to AC; so, it too should be rejected. In short, mixed or half-way positions are false. That leaves two views. On one, all k-conditions are closed and Closure is true; on the other, one (or more) k-condition is not closed and Closure is false.

A second implication of Equivalence applies to some pro-Closure epistemologists, namely those who are sufficiently confident in the truth of Closure that they will reject a theory of knowledge on the grounds that, on that theory, knowledge does not come out closed. Though Equivalence does not tell us whether this is defensible, it does have this interesting implication: via AC, pro-

Closure people of this persuasion can infer that being closed is also an adequacy condition on something's being a k-condition.

There is a related, less benign, implication for those in the pro-Closure camp. This is an implication that *all* proponents of Closure must face, whether they are confident enough to impose Closure as an adequacy condition, or not. It follows from Equivalence in conjunction with an important upshot of our earlier survey. According to the survey, many familiar k-conditions, both internalist and externalist in nature, are *not* closed. This puts pressure on the pro-Closure epistemologist, since it means that she must either reject all of the k-conditions that we surveyed (except the undefeated justification condition, since it is closed), or for any such condition that she incorporates into her analysis of knowledge, she must modify it so that it comes out closed. Either way, she may run into problems. If she takes the first route and opts for an analysis of knowledge that does not include any of the surveyed k-conditions (again, save the undefeated justification condition), it is unclear what her analysis will include. After all, the survey covered *all* of the familiar epistemic conditions in their basic forms. On the other hand, if she takes the second route, as her modifications get more intricate, unless she provides independent support for those modifications, getting by without Closure will appear simpler, less ad hoc, and more attractive.

Still, these are just pressures. Defenders of Closure might overcome them. Suppose, for a moment, that they do. In that case, could anything here help decide

between the pro-Closure position that respects Equivalence and the anti-Closure position that respects Equivalence?

For a certain audience, the answer is ‘yes.’ These are people that subscribe to two claims. They think that the familiar considerations on both sides of the Closure debate are equally compelling; and, in addition, they are sufficiently confident that some particular analysis of knowledge is correct. All these people need to do to overcome the perceived stalemate is let their favored k-conditions dictate a verdict regarding Closure. If all those conditions are closed, then via PC they can infer Closure. On the other hand, if one (or more) of their favored k-conditions is not closed, then via AC they can reject Closure.

Not too many epistemologists seem to fit this profile, though. Most report to having a fairly strong opinion about Closure, almost invariably a pro-Closure opinion. The dialectical situation between those strongly opinionated in favor of Closure and those strongly opinionated against Closure is probably one in which neither side is going to be able to use Equivalence to help win the other side over. For consider how this would be attempted. Those who reject Closure might argue for a certain analysis of knowledge, show that one (or more) condition in their analysis is not closed, then use AC to claim that Closure is false. Meanwhile, those on the pro-Closure side might argue for a certain analysis of knowledge, show that each condition in their analysis is closed, then use PC to derive Closure. Perhaps, neither of these parties should think they are going to win the other side over, though. The main obstacle is the fact that there seems to be, if anything, more

controversy over what the correct k-conditions are than there is over Closure. For this reason, an appeal to some set of k-conditions will probably not be effective in winning over someone sufficiently confident who holds the opposing view about Closure.

Still, Equivalence has important implications: it narrows to two the defensible ways in which analyses of knowledge can be combined with attitudes to Closure; and in tandem with the results of the survey, it puts serious pressure on the pro-Closure outlook that has become orthodoxy.²³

NOTES

¹ This schema says that *C* is closed *under C-entailment*. A distinct schema says that *C* is closed *under known entailment*. The difference between the two is this: where *S meets C with respect to P entails Q* appears in the antecedent of the first, the second has, instead, *S knows that P entails Q*. The arguments that I will give go through for both schemata.

² See Warfield (2004), Brueckner (2004), and Huemer (2005).

³ Pryor (2000) argues that his principle is crucial to the skeptical argument that I outline at the beginning of the next section.

⁴ Transmission was first discussed in Wright (1985). Wright has continued to work on the nature of transmission, especially as it relates to Michael McKinsey's argument for the incompatibility of privileged access and semantic externalism. For that important implication, see Wright (2003) and McKinsey (2003).

⁵ Closure is weaker than the first because Closure does not require that *S* antecedently know that *Q* if *S* is going to know that *P* and know that *P entails Q*. Closure is weaker than Transmission because Closure does not say that if *S* knows that *P* and *S* knows that *P entails Q*, then *S* can come to know that *Q for the very first time*.

⁶ Warfield (2004) goes against this conjecture. From the fact that some k-condition is not closed, he claim, it does not follow that knowledge is not closed. This is because knowledge might still be closed in virtue of the fact that some other k-condition, or combination of k-conditions, is closed. Warfield's point is right as far as it goes – from *just* the fact that a k-condition is not closed, it does not follow that knowledge is not closed. Still, those who wonder how a k-condition could fail to be closed, while knowledge nonetheless is closed, might want this point illuminated – this section is intended to provide the illumination. Ultimately, however, I will argue that Warfield's claim is false. I will argue that on any account of knowledge that includes a familiar k-condition that turns out not to be closed, Closure is false.

⁷ From here on out, I will not mention the latter two claims in Closure's antecedent. They read: "S goes on to infer from these to believe *Q*" and "S thereby bases the belief that *Q* on these beliefs." I offer a rationale for this simplifying assumption at the beginning of Section 4. For the need for such claims, as well as other touch-ups to Closure, see Brueckner (1985), Hales (1995), and Hawthorne (2005).

⁸ Brueckner (2004) p. 334. Brueckner would say the same regarding *P entails Q*, namely if the conditions for knowledge are independent of each other, then S meets D with respect to *P entails Q* does not entail that S meets C with respect to *Q*.

⁹ *If* the survey did turn up a pair of k-conditions that disconfirm the Containment Principle, then the principle would be false. In turn, AC would be false. But this ignores a complication that I will keep mostly suppressed. The complication is owed to the fact that what we are really discussing here are *candidate* k-conditions. Presumably, just a subset of these candidate k-conditions are among the *actual* conditions on knowledge. However, since the following survey will fail to uncover even any candidate k-conditions that count against the Containment Principle, debates over what the correct k-conditions are can be sidestepped.

¹⁰ Fred's belief in the entailment is sensitive if and only if: if the entailment were false, then Fred would not believe it. This is a counterfactual conditional with an impossible antecedent. On the view that there are only possible worlds, and no impossible worlds, counterfactuals of this kind are true. Here I will simply assume this view, largely in order to get on with determining whether sensitivity is closed. Most epistemologists won't find this objectionable, as it is standard, at least in the epistemology literature, to think that the sensitivity condition is always met with respect to necessary truths. See Lewis (1973) p. 24-26 for discussion of counterfactuals with impossible antecedents.

¹¹ See Collins, Hall, and Paul (2004) for a discussion of the main kinds of counterexamples to the counterfactual theory.

¹² Ernest Sosa is the chief proponent of safety. See Sosa (2002).

¹³ Since the process responsible for the belief that *Q* is a belief-dependent process, (roughly) its reliability is determined by the frequency with which it takes one from true beliefs as inputs to true beliefs as outputs versus the frequency with which it takes one from true beliefs as inputs to false belief as outputs. See Goldman (1986), chapter 5.

¹⁴ For this reason, the student meets the basing requirement: his belief that *Q* is based on his belief that *P* and his belief that *P* entails *Q*.

¹⁵ From the point of view of process reliabilism, the fact that the contents of these beliefs make for a deductively valid argument is irrelevant.

¹⁶ See Lehrer and Paxson (1969).

¹⁷ In the case given, the proof for *not-Q* is an overriding defeater for *Q*, since the proof is a reason to believe the denial of *Q*. There is another type of defeater, an undermining defeater – it consists in reasons to distrust how one arrived at the (defeated) belief. So, what about underminers? Does an underminer for one's belief that *Q* entail that either the warrant for one's belief that *P* is defeated or the warrant for one's belief that *P* entails *Q* is defeated? Since the underminer for one's belief that *Q* can take aim at any epistemically essential element in one's arrival at *Q*, among underminers are reasons to think that the belief that *Q* is not based on the beliefs that *P* and *P* entails *Q*. Such an underminer would not defeat the warrant for the belief that *P* or for the warrant for the belief that *P* entails *Q*. It follows that a no undermining defeater condition is not closed. Still, a no overriding defeater condition is closed.

¹⁸ Though I have argued that undefeated justification is closed, we must substitute closed, as well as non-closed, k-conditions into the D slot.

¹⁹ My belief that Larry is married to Pat is sensitive because if Pat had rejected Larry's proposal, Larry would have told me that he is married to someone else; in which case, I would not have believed that he is married to Pat. My belief in the conditional is sensitive since the conditional is necessarily true.

²⁰ Though, detailed examination of several issues is needed to get a handle on exactly how inductively strong the preceding case for AC is. I will just mention one. It concerns how my (admittedly) exclusive focus on basic incarnations of familiar k-conditions bears on the strength of my case. This might severely weaken my case given the fact that proponents of these conditions usually end up defending distinct, more nuanced descendents of these basic conditions. To this, I have two responses. First, since the nuances are not introduced to make the given conditions closed, but rather to achieve other ends like evading counterexamples to a proffered analysis of knowledge, it would be a fluke if the nuances made the resulting conditions closed. Second, the same point also applies in relation to the role played by the Containment Principle: since the nuances are not introduced with anything about the Containment Principle in mind, it would be a sheer fluke if some combination of nuanced conditions violated the Containment Principle.

²¹ See Dretske (1970), Nozick (1981), Goldman (1986), McGinn (1984), and Audi (1988).

²² See Warfield (2004).

²³ Thanks to Albert Casullo, Murali Ramachandran, Jason Stanley, and two referees for comments.

REFERENCES

- Audi, R.: 1988, *Belief, Justification, and Knowledge*, Wadsworth Publishing, Belmont, CA.
- Brueckner, A.: 1985, 'Skepticism and Epistemic Closure,' *Philosophical Topics* 13, 89-118.
- Brueckner, A.: 2004. 'Strategies for refuting closure,' *Analysis* 64, 333-335.
- Collins, J., Hall, N., and Paul, L.: 2004, 'Counterfactuals and Causation: History, Problems, and Prospects,' in Collins, J., Hall, N., and Paul, L (ed.s), *Causation and Counterfactuals*, MIT Press, Cambridge, MA, 1-57.
- Dretske, F.: 1970, 'Epistemic Operators,' *Journal of Philosophy* 67, 1007-1023.
- Dretske, F.: 2005, 'Reply to Hawthorne,' in Steup M. and Sosa E. (ed.s), *Contemporary Debates in Epistemology*, Blackwell Publishing, Malden, MA, 43-46.
- Goldman, A.: 1986, *Epistemology and Cognition*, Harvard University Press, Cambridge, MA.
- Hales, S.: 1995, 'Epistemic Closure Principles,' *Southern Journal of Philosophy* 33, 185-201.
- Hawthorne, J.: 2005, 'The Case for Closure,' in Steup M. and Sosa E. (ed.s), *Contemporary Debates in Epistemology*, Blackwell Publishing, Malden, MA, 26-43.
- Huemer, M.: 2005, 'Logical Properties of Warrant,' *Philosophical Studies* 114, 171-182.
- Lehrer, K. and Paxson, T: 1969, 'Knowledge: Undefeated Justified True Belief,' *Journal of Philosophy* 66, 225-237.
- Lewis, D.: 1973, *Counterfactuals*, Harvard University Press, Cambridge, MA.
- McGinn, C.: 1984, 'The Concept of Knowledge,' *Midwest Studies in Philosophy* IX, 529-554.

McKinsey, M.: 2003, 'Transmission of Warrant and Closure of Apriority,' in Nuccetelli S. (ed.), *New Essays on Semantic Externalism and Self-Knowledge*, MIT Press, Cambridge, MA, 97-115.

Nozick, R.: 1981, *Philosophical Explanations*, Harvard University Press, Cambridge, MA.

Pryor, J.: 2000, 'The Skeptic and the Dogmatist,' *Nous* 34, 517-549.

Sosa, E.: 2002, 'Tracking, Competence, and Knowledge,' in Moser P. (ed.), *The Oxford Handbook of Epistemology*, Oxford University Press, New York, 264-286.

Warfield, T.: 2004, 'When epistemic closure does and does not fail: a lesson from the history of epistemology,' *Analysis* 64, 35-41.

Wright, C.: 1985, 'Facts and Certainty,' *Proceedings of the British Academy* 71, 429-471.

Wright, C.: 2003, 'Some Reflections on the Acquisition of Warrant by Inference,' in Nuccetelli S. (ed.), *New Essays on Semantic Externalism and Self-Knowledge*, MIT Press, Cambridge, MA, 57-77.